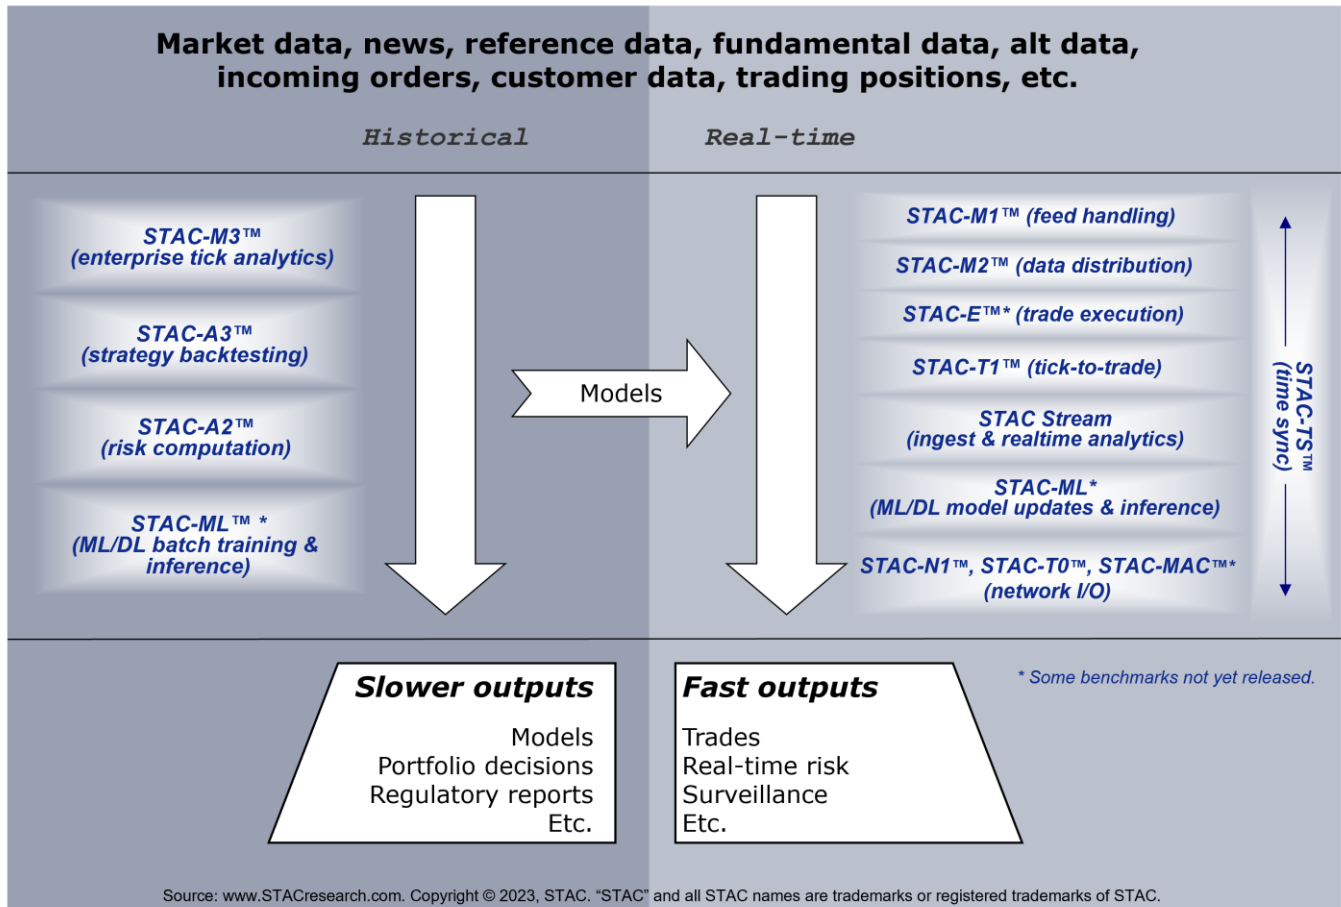




## Capital markets workload categorization

As the figure below shows, STAC Benchmarks™ for capital markets cover key workloads in which financial institutions turn historical and realtime information into action and knowledge. These workloads vary in their intensity of computation, storage I/O, and network I/O. The benchmark suites are categorized primarily as market data (STAC-M™), analytics (STAC-A™), machine learning (STAC-ML™), execution (STAC-E™), "tick-to-trade" (STAC-T™), and time sync (STAC-TS™) with domains that are roughly based on the way that user firms think about their architectures—that is, how they group functional requirements today, divide their technical organizations, and buy products.

# STAC® domains in capital markets



The following tables summarize each benchmark domain. Domains are grouped by whether they focus on a "Fast Workload" (relatively simple computation, with speed of communication being paramount) or a "Big Workload" (where the size of the dataset that must be moved or analyzed, or the magnitude of computation is the biggest challenge—i.e., "big data" or "big compute"). For details or to read STAC Reports™ from the domain, click its title to visit its page on the STAC® site.

## Big Workloads

Domain	Caption	Workload summary	Example metrics	Example products that can be tested
<a href="#">STAC-M3</a>	Tick data analytics	Applying analytics to market data time series. Use cases range from execution reporting and risk management to strategy development. Historical and real-time processing and analytics.	<ul style="list-style-type: none"> <li>• Query-response time</li> <li>• Application-level I/O throughput</li> <li>• Storage efficiency</li> <li>• Data availability</li> </ul>	<ul style="list-style-type: none"> <li>• Tick databases</li> <li>• Storage systems</li> <li>• File systems</li> <li>• Cloud platforms</li> </ul>
<a href="#">STAC-A3</a>	Strategy backtesting	Testing parameterized trading algorithms against historical data.	<ul style="list-style-type: none"> <li>• Speed: time to compute a fixed workload</li> <li>• Capacity: max workload computed in a fixed time</li> <li>• Quality: adherence of results to theoretical values</li> <li>• Efficiency: output per unit of power/space/money</li> </ul>	<ul style="list-style-type: none"> <li>• Storage systems</li> <li>• File systems</li> <li>• Database software</li> <li>• Cloud platforms</li> </ul>
<a href="#">STAC-A2</a>	Risk computations	Computing risk metrics for options, viewed as an end-to-end process and also broken down into discrete compute steps.	<ul style="list-style-type: none"> <li>• Speed: time to compute a fixed workload</li> <li>• Capacity: max workload computed in a fixed time</li> <li>• Quality: adherence of results to theoretical values</li> <li>• Efficiency: output per unit of power/space/money</li> </ul>	<ul style="list-style-type: none"> <li>• CPUs, GPUs, FPGAs</li> <li>• Programming tools</li> <li>• Analytic libraries</li> <li>• Grid software</li> <li>• Clouds</li> </ul>
<a href="#">STAC-ML</a> (batch training and inference subset)	Artificial intelligence ( <i>some not yet released</i> )	Batch-mode inference tasks, as well as offline training of machine learning and deep learning workloads.	<ul style="list-style-type: none"> <li>• Speed: latency of inference</li> <li>• Quality: "goodness" of results</li> <li>• Time to market: cycle time for model training</li> <li>• Efficiency: price performance in the public cloud or output per unit of on-prem resource</li> </ul>	<ul style="list-style-type: none"> <li>• ML/DL frameworks</li> <li>• CPUs, GPUs, FPGAs, custom processors</li> <li>• Inference engines</li> <li>• Libraries</li> <li>• Scaling frameworks</li> <li>• Cloud: IaaS, MLaaS</li> <li>• Storage infrastructure</li> <li>• Techniques</li> </ul>

## Fast Workloads

Domain	Caption	Workload summary	Example metrics	Example products that can be tested
<a href="#">STAC-T0</a>	Network I/O for trading (black-box approach)	Transmitting messages via sockets. Messages follow a tick-to-trade pattern.	<ul style="list-style-type: none"> <li>• “Actionable” latency</li> <li>• FIFO/LIFO/FILO latency</li> </ul>	<ul style="list-style-type: none"> <li>• FPGA platforms and IP cores</li> <li>• Software solutions</li> <li>• NICs/servers/switches</li> </ul>
<a href="#">STAC-T1</a>	Tick to trade	Wire-to-wire I/O performance of a trading system including market data parsing and order generation and transmission.	<ul style="list-style-type: none"> <li>• Market data to order latency at various multiples of market rates</li> </ul>	<ul style="list-style-type: none"> <li>• Feed handling and market gateway software</li> <li>• Server/OS/networking</li> </ul>
<a href="#">STAC-N1</a>	Network I/O for trading (software-based)	Transmitting messages directly via a network API (sockets, RDMA, other). Messages follow a market data pattern.	<ul style="list-style-type: none"> <li>• Per-message API-to-API latency (half round trip)</li> <li>• Max throughput</li> <li>• CPU and memory utilization of sender and receiver</li> </ul>	<ul style="list-style-type: none"> <li>• Network cards and drivers (e.g., kernel bypass)</li> <li>• Server/OS/switches</li> </ul>
<a href="#">“STAC MAC”</a>	MAC/PHY logic ( <i>not yet released</i> )	Handling just the Ethernet MAC & PHY functions	<ul style="list-style-type: none"> <li>• Latency</li> </ul>	<ul style="list-style-type: none"> <li>• IP cores for FPGA</li> </ul>
Under consideration	Network mux ( <i>not yet released</i> )	Multiplexing TCP messages such as orders	<ul style="list-style-type: none"> <li>• Latency</li> </ul>	<ul style="list-style-type: none"> <li>• Multiplexing switches</li> </ul>
<a href="#">STAC-TS</a>	Time synchronization	<p>Synchronizing time, timestamping events, capturing events.</p> <p>Reporting traceability of timestamping points in an enterprise to a standard (e.g., UTC, NIST).</p>	<ul style="list-style-type: none"> <li>• Accuracy of application timestamps relative to an external reference</li> <li>• Accuracy of application timestamps relative to the host clock</li> <li>• Accuracy of master clocks under exception conditions</li> <li>• Accuracy of network timestamps</li> <li>• Capacity of network capture</li> </ul>	<ul style="list-style-type: none"> <li>• Timestamping API calls</li> <li>• Master clocks</li> <li>• PTP solutions</li> <li>• NTP solutions</li> <li>• Switches</li> <li>• NICs</li> <li>• Capture cards</li> <li>• Aggregation devices</li> </ul>
<a href="#">STAC-M1</a>	Direct feed integration	Taking inbound market data messages from exchanges, normalizing and caching them, and making them available via an API.	<ul style="list-style-type: none"> <li>• Latency from exchange message hitting the wire to normalized update coming through API</li> <li>• Max throughput, and latency at that throughput</li> <li>• CPU/memory consumption of client</li> </ul>	<ul style="list-style-type: none"> <li>• Ticker plant software</li> <li>• Ticker plant appliances</li> <li>• Acceleration cards</li> <li>• Server/OS/networking</li> </ul>
<a href="#">STAC-M2</a>	Market data distribution	Taking normalized market data streams through a publisher API and delivering them to multiple consumers via a subscriber API.	<ul style="list-style-type: none"> <li>• Latency from the moment a message is ready for distribution to the moment it exits the subscriber API.</li> <li>• Latency to “undisturbed consumers” when other consumers are starved of resources</li> </ul>	<ul style="list-style-type: none"> <li>• Messaging software</li> <li>• Messaging appliances</li> <li>• Market data platforms</li> <li>• Server/OS/networking</li> </ul>

<a href="#">STAC-E</a>	Trade-execution ( <i>not yet released</i> )	Performing functions on trade messages inbound via APIs and/or wire protocols and transmitting them to a destination. Functions may include validation, pre-trade risk and compliance checks, order-state management, matching, and smart order routing. Includes internal messaging via persistent methods (resilience to internal failures).	<ul style="list-style-type: none"> <li>• One-way latency according to message type</li> <li>• Two way latency (e.g., order to fill, cancel to ACK, order to quote)</li> <li>• Internal API-to-API latency of persistent messaging</li> <li>• Max throughput</li> </ul>	<ul style="list-style-type: none"> <li>• Market- or client-facing execution gateways</li> <li>• Messaging software, messaging appliances</li> <li>• OMS</li> <li>• Smart order routers</li> <li>• Matching engines</li> <li>• Server/OS/networking</li> </ul>
<a href="#">STAC-M3</a> "Jalua"	Streaming ingest and analytics	Various use cases that require ingest of streaming data and either event-driven or on-demand analytics.	<ul style="list-style-type: none"> <li>• Capacity: Max ingest rate</li> <li>• Latency: time for inbound data to become available for analytics or to result in outbound notifications</li> <li>• Efficiency: work per unit of power/space/dollar</li> </ul>	<ul style="list-style-type: none"> <li>• Data management software</li> <li>• Memory, storage, interconnects</li> <li>• Deployed systems and clouds</li> </ul>
<a href="#">STAC-ML</a> (model updates and inference subset)	Artificial intelligence (some <i>not yet released</i> )	Batch-mode inference tasks, as well as offline training of machine learning and deep learning workloads.	<ul style="list-style-type: none"> <li>• Speed: latency of inference</li> <li>• Quality: "goodness" of results</li> <li>• Time to market: cycle time for model training</li> <li>• Efficiency: price performance in the public cloud or output per unit of on-prem resource</li> </ul>	<ul style="list-style-type: none"> <li>• ML/DL frameworks</li> <li>• CPUs, GPUs, FPGAs, custom processors</li> <li>• Inference engines</li> <li>• Libraries</li> <li>• Scaling frameworks</li> <li>• Cloud: IaaS, MLaaS</li> <li>• Storage infrastructure</li> <li>• Techniques</li> </ul>